

There are Plane Spanners of Maximum Degree 4

Nicolas Bonichon* Iyad Kanj† Ljubomir Perković† Ge Xia‡

Abstract

Let \mathcal{E} be the complete Euclidean graph on a set of points embedded in the plane. Given a constant $t \geq 1$, a spanning subgraph G of \mathcal{E} is said to be a t -*spanner*, or simply a spanner, if for any pair of vertices u, v in \mathcal{E} the distance between u and v in G is at most t times their distance in \mathcal{E} . A spanner is *plane* if its edges do not cross.

This paper considers the question: “What is the smallest *maximum degree* that can always be achieved for a *plane* spanner of \mathcal{E} ?” Without the planarity constraint, it is known that the answer is 3 which is thus the best known lower bound on the degree of any plane spanner. With the planarity requirement, the best known upper bound on the maximum degree is 6, the last in a long sequence of results improving the upper bound. In this paper we show that the complete Euclidean graph always contains a plane spanner of maximum degree at most 4 and make a big step toward closing the question. Our construction leads to an efficient algorithm for obtaining the spanner from Chew’s L_1 -Delaunay triangulation.

*LaBRI, Université Bordeaux 1. bonichon@labri.fr. This work was partially supported by ANR grant JCJC EGOS ANR-12-JS02-002-01.

†School of Computing, DePaul University. {ikanj,lperkovic}@cs.depaul.edu.

‡Department of Computer Science, Lafayette College. xiag@lafayette.edu.

1 Introduction

Let \mathcal{E} be the complete Euclidean graph on a set of points P embedded in the plane. Given a constant $t \geq 1$, a spanning subgraph G of \mathcal{E} is said to be a t -*spanner*, or simply a *spanner*, if for any pair of vertices u, v in \mathcal{E} the distance between u and v in G is at most t times their distance in \mathcal{E} . The constant t is referred to as the *stretch factor*. A spanner is *plane* if its edges do not cross.

In this paper, we consider the following question: *What is the smallest maximum degree that can always be achieved for plane spanners of complete Euclidean graphs?* Or, to put it more precisely: *What is the smallest d such that for some constant $t \geq 1$ there always exists a plane t -spanner of maximum degree at most d on any set of points on the plane?* This fundamental question was raised by Bose and Smid [BS13] in their recent survey of geometric problems. It is a natural extension to classical questions on spanners of complete Euclidean graphs, and Delaunay triangulations in particular.

In the mid-1980s, the fundamental question of whether a plane spanner of \mathcal{E} always exists was considered. In his seminal 1986 paper, Chew answered the question in the affirmative [Che86]. He proved, in particular, that the L_1 -Delaunay triangulation of P , i.e. the dual of the Voronoi diagram of P based on the L_1 -distance, is a $\sqrt{10}$ -spanner of \mathcal{E} . Chew’s result was followed by a series of papers demonstrating that other Delaunay triangulations are plane spanners as well. In 1987, Dobkin *et al.* [DFS90] were successful in showing that the (classical) L_2 -Delaunay triangulation of P , i.e. the dual of the Voronoi diagram of P based on the L_2 -distance (i.e., the Euclidean distance) is a spanner as well. The bound on the stretch factor they obtained was subsequently improved by Keil and Gutwin [KG92] as shown in Table 1. In the meantime, Chew [Che89] showed that the *TD*-Delaunay triangulation—again a dual of a Voronoi diagram but this time defined using a distance function based on an equilateral triangle rather than a square (L_1 -distance) or a circle (L_2 -distance)—is a 2-spanner.

The bound on the stretch factor of an L_2 -Delaunay triangulation by Keil and Gutwin stood unchallenged for many years until Xia recently improved the bound to below 2 [Xia13] (see Table 1). Recently as well, Bonichon *et al.* [BGHP12] improved Chew’s original bound on the stretch factor of the L_1 -Delaunay triangulation to $\sqrt{4 + 2\sqrt{2}}$ and showed this bound to be tight.

Minimizing the stretch factor of a plane spanner of \mathcal{E} is one natural goal. Another one is minimizing the maximum degree of the plane spanner. This restriction eliminates, for example, the various Delaunay triangulations because they can have unbounded degree. The lower bound on the maximum degree of a spanner is 3, because a Hamiltonian path through a set of n points arranged in a grid has stretch factor $\Omega(\sqrt{n})$. Work on bounded degree *but not necessarily plane* spanners of \mathcal{E} closely followed the above-mentioned work on plane spanners. In a 1992 breakthrough, Salowe

Paper	Spanner	Stretch factor bound
Chew [Che86]	L_1 -Delaunay	$\sqrt{10} \approx 3.16$
Bonichon et al. [BGHP12]	L_1 -Delaunay	$\sqrt{4 + 2\sqrt{2}} \approx \mathbf{2.61}$
Dobkin <i>et al.</i> [DFS90]	L_2 -Delaunay	$\frac{\pi(1+\sqrt{5})}{2} \approx 5.08$
Keil & Gutwin [KG92]	L_2 -Delaunay	$\frac{4\pi}{3\sqrt{3}} \approx 2.42$
[Xia13]	L_2 -Delaunay	1.998
Chew [Che89]	<i>TD</i> -Delaunay	$\mathbf{2}$

Table 1: Key results on (unbounded degree) plane spanners; tight bounds are in bold.

Paper	Δ	Stretch factor bound
Bose <i>et al.</i> [BGS05]	27	$(\pi + 1)C_0 \approx 8.27$
Li and Wang [LW04]	23	$(1 + \pi \sin \frac{\pi}{4})C_0 \approx 6.43$
Bose <i>et al.</i> [BSX09]	17	$(2 + 2\sqrt{3} + \frac{3\pi}{2} + 2\pi \sin(\frac{\pi}{12}))C_0 \approx 23.56$
Kanj and Perković [KP08]	14	$(1 + \frac{2\pi}{14 \cos(\frac{\pi}{14})})C_0 \approx 2.91$
Bonichon <i>et al.</i> [BGHP10]	6	6
Bose <i>et al.</i> [BCCY12]	6	$1/(1 - \tan(\pi/7)(1 + 1/\cos(\pi/14)))C_0 \approx 81.66$
This paper	4	$\sqrt{4 + 2\sqrt{2}}(1 + \sqrt{2})^2(3 + \sqrt{2})^6 \approx 112676$

Table 2: Results on plane spanners with maximum degree bounded by Δ . The constant $C_0 = 1.998$ is the best known upper bound on the stretch factor of the L_2 -Delaunay triangulation [Xia13]. The stretch factor bound in this paper can be made much tighter with a more careful analysis.

[Sal94] proved the existence of spanners of maximum degree at most 4. The question was then resolved by Das and Heffernan [DH96] who showed that spanners of maximum degree at most 3 always exist.

The focus in this line of research was to prove the existence of low degree spanners and the techniques developed to do so were not tuned towards constructing spanners that had both low degree *and* low stretch factor. Furthermore, the bounded-degree spanners shown to exist were not guaranteed to be plane. In recent years, bounded degree plane spanners have been used as the building block of wireless network topologies. Emerging wireless distributed system technologies, such as wireless ad-hoc and sensor networks, are often modeled as proximity graphs in the Euclidean plane. Spanners of proximity graphs represent topologies that can be used for efficient unicasting, multicasting, *and/or* broadcasting. For these applications, in addition to low stretch factor, spanners are typically required to be plane and have bounded degree. The planarity requirement is for efficient routing (see [BMSU01]), while the bounded degree requirement is due to the physical limitations of wireless devices (see [WL06]).

Bose *et al.* [BGS05] were the first to show how to extract a spanning subgraph of the classical L_2 -Delaunay triangulation that is a bounded-degree, plane spanner of \mathcal{E} . The maximum degree and stretch factor bounds they obtained were subsequently improved by Li and Wang [LW04], by Bose *et al.* [BSX09], and by Kanj and Perković [KP08] (see all bounds in Table 1). The approach used in all these results was to extract a bounded degree spanning subgraph of the classical L_2 -Delaunay triangulation and the main goal was to obtain a bounded-degree plane spanner of \mathcal{E} with the smallest possible stretch factor.

Recently, Bonichon *et al.* [BGHP10] focused on lowering the bound on the maximum degree of a plane spanner and developed a new approach. Instead of using the classical L_2 -Delaunay triangulation as the starting point of the spanner construction, they used the TD -Delaunay triangulation defined by Chew [Che89]. They achieved a significant decrease in the bound on the maximum degree: from 14 down to 6. The plane spanner they constructed also had a surprisingly small stretch factor of 6. Independently, Bose *et al.* [BCCY12] have also been able to obtain a plane spanner of maximum degree at most 6, by starting from the L_2 -Delaunay triangulation; the spanner they obtain has the additional property of being *strong* which means that between every pair of vertices u and v there is, in the spanner, a path that consists of edges whose length is no more than the Euclidean distance between u and v .

In this paper, we push the bound on the maximum degree of a plane spanner from 6 down

to 4 and make a big step toward closing a fundamental question. Interestingly, the starting point for our spanner construction is Chew’s original L_1 -Delaunay triangulation, a graph that has been largely overlooked in the last quarter century. We define this triangulation, and the equivalent L_∞ -Delaunay triangulation, in the next section. In Section 3, we introduce a key tool: a directed version of the L_∞ -distance-based Yao graph Y_4^∞ introduced by Bose *et al.* [BDD⁺12]. *En passant*, we prove that Y_4^∞ is a plane $\sqrt{20 + 14\sqrt{2}} \approx 6.3$ -spanner of \mathcal{E} . Then, in Section 4, we define *standard paths* between the endpoints of every edge in Y_4^∞ . In Section 5, we construct a subgraph H_8 of Y_4^∞ of maximum degree at most 8 and show that it is a spanner by proving that it contains short standard paths. Finally, in Section 6, we show that some edges in H_8 are redundant and we remove them, while adding new shortcut edges, to obtain H_4 , a spanner of maximum degree at most 4. While the proofs in the paper are quite technical, our construction leads to a simple and efficient algorithm for computing the spanner.

2 Preliminaries

Let P be a set of points in the two-dimensional Euclidean space. The Euclidean graph \mathcal{E} of P is the complete weighted graph embedded in the plane whose nodes are identified with the points of P . We assume that a coordinate system is associated with the Euclidean plane and thus every point can be specified by its x and y coordinates. For every pair of nodes u and w , we identify edge (u, w) with the straight line segment $[uw]$ and associate an edge length equal to the Euclidean distance $d_2(u, w) = \sqrt{d_x(u, w)^2 + d_y(u, w)^2}$ where $d_x(u, w)$ (resp. $d_y(u, w)$) is the difference between the x (resp. y) coordinates of u and w . Given a constant $t \geq 1$, we say that a subgraph H of a graph G is a t -spanner, or simply a *spanner*, of G if for any pair of vertices u, v of G , the distance between u and v in H is at most t times the distance between u and v in G ; the constant t is referred to as the *stretch factor* of H (with respect to G). We will say that H is a t -spanner, or simply a spanner, if it is a t -spanner of \mathcal{E} .

In the introduction we defined the L_1 -Delaunay triangulation as the dual of the Voronoi diagram based on the L_1 -distance defined as $d_1(u, w) = d_x(u, w) + d_y(u, w)$ for two points u and w . In this paper, our working definition is an alternate but equivalent one. Let a *square* in the plane be a square whose sides are parallel to the x and y axes and let a *tipped square* be a square tipped at 45° . For every pair of points $u, v \in P$, (u, v) is an edge in the L_1 -Delaunay triangulation of P iff there is a tipped square that has u and v on its boundary and is *empty* (i.e., it contains no point of P in its interior). The assumption in this definition is that the points of P are in *general position* which implies that no four points lie on the boundary of a tipped square. With this assumption, an L_1 -Delaunay triangulation is indeed a plane graph whose interior faces are all triangles.

If a square with sides parallel to the x and y axes, rather than a tipped square, is used in the above definition then a different triangulation is obtained; it corresponds to the dual of the Voronoi diagram based on the L_∞ -distance $d_\infty(u, w) = \max\{d_x(u, w), d_y(u, w)\}$. Here again the assumption is that points of P are in *general position* which in *this* case implies that no four points lie on the boundary of a square. We refer to the resulting triangulation as the L_∞ -Delaunay triangulation. This triangulation is nothing more than the L_1 -triangulation of the set of points P after rotating all the points by 45° around the origin. Therefore Chew’s bound of $\sqrt{10}$ on the stretch factor of the L_1 -Delaunay triangulation ([Che86]) applies to L_∞ -Delaunay triangulations as well. In the remainder of this paper, we will be using L_∞ -Delaunay (rather than L_1 -) triangulations because we will be (mostly) using the L_∞ -distance, and squares rather than tipped squares.

In order to avoid technical difficulties we make the usual assumption that points of P are in

general position which for us means that 1) no four points lie on the boundary of a square and 2) no two points have the same x or y coordinate. Note that it is always possible to perturb the points slightly so they end up in general position and so that a plane spanner on the perturbed points corresponds to a plane spanner on the original points. Therefore, the main result in this paper holds for all sets of points and not just for points in general position.

3 A Yao subgraph of the L_∞ -Delaunay triangulation

In this section we describe the first step in the construction of our spanner of \mathcal{E} , the complete Euclidean graph on a set of points P . The result of the first step is a version of the Yao subgraph of \mathcal{E} on four cones and first defined by Bose *et al.* [BDD⁺12].

A *cone* is the open region in the plane between two rays that emanate from the same point. With every point u of P we associate four disjoint 90° cones emanating from u : they are defined by the translation of the x - and y -axis from the origin to point u and exclude the translated axes. We label the cones 0, 1, 2, and 3, in counter-clockwise order, starting with the cone corresponding to the first quadrant. Given a cone i , the counter-clockwise next cone is cone $i + 1$, whereas the clockwise next cone is cone $i - 1$; we assume that arithmetic on the labels is done modulo 4 so that cone $i + 1$ and cone $i - 1$ are well defined. Our general position assumption ensures that no point lies on the boundary of another point's cone.

Given two points v and w , we define $R(v, w)$ to be the rectangle, with sides parallel to the x and y axes, having v and w as vertices. The rectangle has positive area because (of our general assumption that) no two points share the same x or y coordinate. For a point v and cone i of v , we denote by $S_v^i(s)$ the $s \times s$ square having v as a vertex and whose two sides match the boundary of cone i of v , and by S_v^i the square $S_v^i(s)$ with the largest s that contains no points of P in its interior (see Figure 1).

The following is the first step of our spanner construction:

Step 1 For every node v of P , we choose in each non-empty cone of v the shortest edge of \mathcal{E} incident to v according to the L_∞ -distance, breaking ties arbitrarily, and we give it an orientation out of v . (See Figure 1.)

We name the resulting directed graph $\overrightarrow{Y_4^\infty}$ and denote an edge of $\overrightarrow{Y_4^\infty}$ from node v to node w using notation $(\overrightarrow{v, w})$ (see Figure 5-(a)). If edge $(\overrightarrow{v, w})$ is in $\overrightarrow{Y_4^\infty}$ then w must lie on the boundary of S_v^i for some cone i . Because for every $(\overrightarrow{v, w}) \in \overrightarrow{Y_4^\infty}$ there is an empty square with v and w on




Figure 1: Definition of S_v^i and orientation of edges in $\overrightarrow{Y_4^\infty}$.




Figure 2: Illustrations for Observation 3.1.

its boundary, (v, w) must be an edge in the L_∞ -Delaunay triangulation T of the points in P (see Figure 5-(a)). Thus the undirected graph obtained by removing the orientations of edges in $\overrightarrow{Y_4^\infty}$ is a subgraph of T which we denote as Y_4^∞ (just as in [BDD⁺12]). For a given edge $(v, w) \in Y_4^\infty$ it is possible that orientation $(\overrightarrow{v, w}) \in \overrightarrow{Y_4^\infty}$, that orientation $(\overrightarrow{w, v}) \in \overrightarrow{Y_4^\infty}$, or that both orientations are in $\overrightarrow{Y_4^\infty}$. We will call (v, w) *uni-directional* in the first two cases and *bi-directional* in the third case.

If an edge (u, v) of Y_4^∞ is in cone i of u then it must also be in cone $i + 2$ of v . One possibility is that (u, v) is the only edge incident to u in its cone i and the only edge incident to v in its cone $i + 2$; we call such an edge a *mutually-single* edge and note that it must be bi-directional. If that is not the case, there must be either two or more edges of Y_4^∞ incident to u in its cone i , or two or more edges of Y_4^∞ incident to v in its cone $i + 2$, or both. We call edge (u, v) *dual* if there are two or more edges of Y_4^∞ incident to u in its cone i and two or more edges of Y_4^∞ incident to v in its cone $i + 2$. Finally, given a node u and cone i of u , we define the *fan* of u in cone i to be the sequence, in counter-clockwise order, of all edges of Y_4^∞ incident to u in its cone i .

Observation 3.1 *For every node u with a fan $(u, v_1), \dots, (u, v_k)$, with $k \geq 2$, in its cone i :*

- (a) $R(u, v_l)$ is empty, for every $l = 1, 2, \dots, k$ (see Figure 2-(a)).
- (b) For every $l \in \{2, \dots, k - 1\}$, edge (u, v_l) is the only edge incident to v_l in its cone $i + 2$ and $(\overrightarrow{v_l, u}) \in \overrightarrow{Y_4^\infty}$ (see Figure 2-(a)).
- (c) For $l \in \{1, k\}$, if edge (u, v_l) is not dual then (u, v_l) is the only edge incident to v_l in its cone $i + 2$ and $(\overrightarrow{v_l, u}) \in \overrightarrow{Y_4^\infty}$.
- (d) For every $l \in \{1, \dots, k - 1\}$, v_l lies in cone $i + 3$ of v_{l+1} , v_{l+1} lies in cone $i + 1$ of v_l , and $R(v_l, v_{l+1})$ is empty (see Figure 2-(b)).
- (e) For every $l \in \{1, \dots, k - 1\}$, **if** $d_1(v_l, u) \leq d_1(v_{l+1}, u)$ (resp. $d_1(v_l, u) \geq d_1(v_{l+1}, u)$) **then** $(\overrightarrow{v_l, v_{l+1}}) \in \overrightarrow{Y_4^\infty}$ (resp. $(\overrightarrow{v_{l+1}, v_l}) \in \overrightarrow{Y_4^\infty}$); furthermore, if edge (v_l, v_{l+1}) is uni-directional then the converse is also true (see Figure 2-(c)).

Proof. Since $(u, v_l) \in Y_4^\infty$, there is an empty square with u and v_l on its boundary. Rectangle $R(u, v_l)$ is contained in this square and thus has no point of P in its interior which proves part (a).

For every $l \in \{2, \dots, k - 1\}$, since $R(u, v_l)$ is empty, any edge other than (u, v_l) incident to v_l in its cone $(i + 2)$ must either intersect edge (u, v_{l-1}) or (u, v_{l+1}) , contradicting the planarity of Y_4^∞ (recall

that Y_4^∞ is a subgraph of the L_∞ -Delaunay triangulation T). Thus, for every $l \in \{2, \dots, k-1\}$, (u, v_l) is the only edge incident to v_l in its cone $i+2$. By construction of $\overrightarrow{Y_4^\infty}$, $(\overrightarrow{v_l, u})$ must be in $\overrightarrow{Y_4^\infty}$ and $S_{v_l}^{i+2}$ must have u on its boundary which proves part (b). The previous statement is also true for $l \in \{1, k\}$ if (u, v_l) is the only edge incident to v_l in its cone $i+2$, which proves part (c).

Since edge (u, v_{l+1}) is counter-clockwise from edge (u, v_l) inside cone i of u , for every $l \in \{1, \dots, k-1\}$, node v_{l+1} cannot be in cone $i+3$ of v_l . Node v_{l+1} cannot be in cone i of v_l because rectangle $R(u, v_{l+1})$ would then contain node v_l . Node v_{l+1} also cannot be in cone $i+2$ of v_l because rectangle $R(u, v_l)$ would then contain point v_{l+1} . So v_{l+1} must be in cone $i+1$ of v_l and thus v_l is in cone $i+3$ of v_{l+1} . If rectangle $R(v_l, v_{l+1})$ contains a point of P for some $l = 1, \dots, k-1$, let w be a point inside $R(v_l, v_{l+1})$ such that rectangle $R(w, u)$ is empty. Because u lies inside cone $i+2$ of w , by construction of $\overrightarrow{Y_4^\infty}$ there must be an edge in $\overrightarrow{Y_4^\infty}$ out of w in its cone $i+2$. Because $R(w, u)$ is empty, any edge incident to w in its cone $i+2$ whose endpoint is not u would have to intersect edge (u, v_l) or (u, v_{l+1}) , contradicting the planarity of Y_4^∞ . Therefore, (u, w) would have to be an edge of Y_4^∞ lying between (u, v_l) and (u, v_{l+1}) in cone i of u which contradicts our assumption and proves part (d).

Let $d_1(v_l, u) \leq d_1(v_{l+1}, u)$ for some $l \in \{1, \dots, k-1\}$ (the case $d_1(v_l, u) \geq d_1(v_{l+1}, u)$ is symmetric). We assume first that u lies on the boundary of $S_{v_{l+1}}^{i+2}$; from the above proof of part (b), that is not the case only if $l = k-1$ and $(u, v_{l+1} = v_k)$ is dual. Since square $S = S_{v_l}^{i+1}(d_\infty(v_l, v_{l+1}))$ lies inside $R(v_l, v_{l+1}) \cup S_{v_{l+1}}^{i+2}$, S must be empty and no point of P other than v_l and v_{l+1} lies on the boundary of S . If $l = k-1$, $(u, v_{l+1} = v_k)$, but u does not lie on the boundary of $S_{v_k}^{i+2}$ then, by construction of $\overrightarrow{Y_4^\infty}$, v_k must lie on the boundary of S_u^i . In that case, $S = S_{v_l}^{i+1}(d_\infty(v_l, v_{l+1}))$ lies inside $R(v_l, v_{l+1}) \cup S_u^i$ and so S is empty and no point of P other than v_l and v_{l+1} lies on the boundary of S . Since v_{l+1} lies on the boundary of S , square $S_{v_l}^{i+1}$ is exactly square S . Since no point of P other than v_l and v_{l+1} lies on the boundary of $S_{v_l}^{i+1}$, $(\overrightarrow{v_l, v_{l+1}}) \in \overrightarrow{Y_4^\infty}$ and (e) follows. \square

Let $(u, v_1), (u, v_2), \dots, (u, v_k)$, with $k \geq 2$, be the fan of u in its cone i . We call (u, v_1) and (u, v_k) the *first* and *last* edge, respectively, in cone i of u .¹ We call any remaining edge (u, v_l) ($1 < l < k$) a *middle edge* of u ; we say that an edge is a *middle edge* if it is a *middle edge* of one of its endpoints. By Observation 3.1-(e), $(v_1, v_2), (v_2, v_3), \dots, (v_{k-1}, v_k)$ are all edges in Y_4^∞ . We call these edges *canonical edges* of u in its cone i and we say that an edge is *canonical* if it is a canonical edge of some node. We make a few observations to differentiate middle, dual, and canonical edges:

Observation 3.2 *Let (u, v) be an edge of Y_4^∞ lying in cone i of u .*

- (a) *If (u, v) is dual then (u, v) is the first edge in cone i of u and cone $i+2$ of v or the last edge in cone i of u and cone $i+2$ of v .*
- (b) *If (u, v) is a uni-directional canonical edge such that $(\overrightarrow{v, u}) \in \overrightarrow{Y_4^\infty}$ then (u, v) is the first or last edge in cone i of u , the only edge in cone $i+2$ of v , and a canonical edge of just one node.*
- (c) *(u, v) can belong to at most one of the following categories: middle, dual, or canonical.*

Proof. By Observation 3.1-(b), if (u, v) is dual then it must be the first or last edge in cone i of u and cone $i+2$ of v . W.l.o.g. we assume that (u, v) is first in cone i of u . Since $R(v, u)$ is empty (Observation 3.1-(a)) and because Y_4^∞ is planar, any edge incident to v in its cone $i+2$ other than (v, u) must be counter-clockwise from (v, u) in cone $i+2$ of v . Thus (v, u) must also be first in cone $i+2$ of v which proves part (a).

¹The first and the last edge in a cone of u are defined only if there are two or more edges incident to u in the cone.




Figure 3: Proof of Observation 3.3.

If (u, v) is a canonical edge of some node w then, by Observation 3.1-(d), w must lie in cones $i - 1$ of u and v or in cones $i + 1$ of u and v . W.l.o.g. we assume the former. If (u, v) is also uni-directional and $(\vec{v}, \vec{u}) \in \overline{Y_4^\infty}$ then there must be at least one more edge of Y_4^∞ incident to u in its cone i . Let $(u, v_1), \dots, (u, v_k)$, with $k \geq 2$, be the fan of u in its cone i . If $v = v_l$ for some $l > 1$ then either v_1 is contained in $R(v, w)$, which contradicts Observation 3.1-(a), or (v, w) intersects (u, v_1) , which contradicts the planarity of Y_4^∞ . Therefore, v must be v_1 . This same argument can be used to show that (u, v) cannot be a canonical edge of a node in cones $i + 1$ of u and v . Thus (u, v) is a canonical edge of node w only. Because, by Observation 3.1-(a), $R(u, w)$ and $R(u, v)$ are empty, any edge incident to v in its cone $i + 2$ other than (u, v) would have to intersect (u, v_2) or (u, w) , which contradicts the planarity of Y_4^∞ . Therefore, (u, v) is the only edge in cone $i + 2$ of v , which completes the proof of part (b).

If (u, v) is a middle edge then, by Observation 3.1-(b), it cannot be the first or the last edge in cone i of u and in cone $i + 2$ of v . This, together with parts (a) and (b), proves part (c). \square

Lemma 3.3 Y_4^∞ is a plane $(1 + \sqrt{2})$ -spanner of the L_∞ -Delaunay triangulation T and also a $(1 + \sqrt{2})\sqrt{4 + 2\sqrt{2}}$ -spanner of \mathcal{E} .

Proof. Let (u, v) be an edge in T that is not in Y_4^∞ and let (u, v) lie in cone i of u . Since $(u, v) \in T$, there exists a square S circumscribing (u, v) and containing no points of P in its interior. W.l.o.g, we assume that v is in cone 0 of u . If u and v lie on adjacent sides of S then either S_u^0 is contained in S and has v on its boundary, implying $(\vec{u}, \vec{v}) \in \overline{Y_4^\infty}$, or S_v^2 is contained in S and has u on its boundary, implying $(\vec{v}, \vec{u}) \in \overline{Y_4^\infty}$, contradicting the assumption that $(u, v) \notin Y_4^\infty$. Therefore we can assume that u and v lie in the interior of opposite sides of the square S , say the bottom and top sides, respectively. We can also assume that S also has another point, w , on its boundary, say on the interior of the right side of S (otherwise we translate S to the right until that occurs) as shown in Figure 3. Because S is devoid of points of P , edges (u, w) and (w, v) are in T . Since $d_x(u, w)$ and $d_y(u, w)$ are both less than $d_y(u, v)$, it follows that $d_\infty(u, w) < d_\infty(u, v)$. If $d_x(u, w) \geq d_y(u, w)$ then square $S_u^0(d_x(u, w))$ is contained inside square S and is thus empty and $(\vec{u}, \vec{w}) \in \overline{Y_4^\infty}$; otherwise square $S_w^2(d_y(u, w))$ is empty and $(\vec{w}, \vec{u}) \in \overline{Y_4^\infty}$. So $(u, w) \in Y_4^\infty$. A similar argument can be used to show that (w, v) is in Y_4^∞ as well. Given that u and v lie on the bottom and top sides of square S and w lies on the right side of S , it follows that $d_2(u, w) + d_2(w, v) \leq (1 + \sqrt{2})d_2(u, v)$.

Since the L_∞ -Delaunay triangulation is a $\sqrt{4 + 2\sqrt{2}}$ -spanner of the complete Euclidean graph [BGHP12], it follows that Y_4^∞ is a $(1 + \sqrt{2})\sqrt{4 + 2\sqrt{2}}$ -spanner of \mathcal{E} . \square

4 Anchors and standard paths

In cone i of some node u , either there is no edge of Y_4^∞ incident to u or there is exactly one edge of $\overrightarrow{Y_4^\infty}$ out of u and any number of edges of $\overrightarrow{Y_4^\infty}$ into u . In this section we describe how, under certain conditions, we choose a special *anchor* among all those edges. We then use anchors to define special paths between endpoints of every edge in Y_4^∞ . We start with some definitions:

Definition 1 Let $(u, v_1), \dots, (u, v_k)$ be the fan of u in its cone i . For any $r, s \in \{1, \dots, k\}$, we define $\text{cpath}_u(v_s, v_r)$ to be the path v_s, v_{s+1}, \dots, v_r (if $s \leq r$) or v_s, v_{s-1}, \dots, v_r (if $s > r$) in Y_4^∞ . We will say that $\text{cpath}_u(v_s, v_r)$ is a *uni-directional canonical path* if every edge in the path is uni-directional and v_s, \dots, v_r forms a directed path from v_s to v_r in $\overrightarrow{Y_4^\infty}$. A uni-directional canonical path *ending at v_r* is *maximal* if it is not contained in any other uni-directional canonical path ending at v_r .

For instance, in Figure 5-(a), (u_9, u_3) and (u_9, u_{20}) belong to the fan of u_9 hence the path $\text{cpath}_{u_9}(u_3, u_{20})$ is well-defined. The path $\text{cpath}_{u_9}(u_{20}, u_{17})$ is a maximal uni-directional canonical path ending at u_{17} . Note that if v_s, v_{s+1}, \dots, v_r is a *maximal uni-directional canonical path ending at v_r* then either $s = 1$ or $(\overrightarrow{v_s, v_{s+1}}) \in \overrightarrow{Y_4^\infty}$. Similarly, if v_s, v_{s-1}, \dots, v_r is a *maximal uni-directional canonical path ending at v_r* then either $s = k$ or $(\overrightarrow{v_s, v_{s+1}}) \in \overrightarrow{Y_4^\infty}$. We can now define anchor edges:

Definition 2 For every node u and every cone i of u containing an edge incident to u :

- (a) If $(u, v) \in Y_4^\infty$ is a mutually-single edge in cone i of u , we define (u, v) to be the anchor chosen by u in cone i .
- (b) If there are two or more edges of Y_4^∞ incident to u in its cone i , let $(u, v_1), \dots, (u, v_k)$, with $k \geq 2$, be the fan of u in cone i and let $(\overrightarrow{v_l, v_l})$, for some $l \in \{1, \dots, k\}$, be the only edge of $\overrightarrow{Y_4^\infty}$ in cone i of u that is outgoing with respect to u :
 - (i) If $l \geq 2$ and $(\overrightarrow{v_{l-1}, v_l}) \in \overrightarrow{Y_4^\infty}$ but $(\overrightarrow{v_l, v_{l-1}}) \notin \overrightarrow{Y_4^\infty}$, let $l' < l$ be such that $\text{cpath}_u(v_{l'}, v_l)$ is a maximal uni-directional canonical path ending at v_l ; we define $(u, v_{l'})$ to be the anchor chosen by u in cone i .
 - (ii) Otherwise, if $l \leq k - 1$ and $(\overrightarrow{v_{l+1}, v_l}) \in \overrightarrow{Y_4^\infty}$ but $(\overrightarrow{v_l, v_{l+1}}) \notin \overrightarrow{Y_4^\infty}$, let $l' > l$ be such that $\text{cpath}_u(v_{l'}, v_l)$ is a maximal uni-directional canonical path ending at v_l ; we define $(u, v_{l'})$ to be the anchor chosen by u in cone i .
 - (iii) Otherwise, we define $(u, v_{l'}) = (u, v_l)$ to be the anchor chosen by u in cone i .

We use the notation $\text{anchor}_i(u)$ to denote the anchor edge chosen by node u in its cone i . In Figures 5-(b)-(d), an anchor $(u, v_{l'})$ is represented by a thick edge with an arrow toward $v_{l'}$ at the end of the edge. In Figure 5-(b), $\text{anchor}_0(u_{22}) = (u_{22}, u_{15})$ illustrates case (a), $\text{anchor}_3(u_9) = (u_9, u_{20})$ illustrates case (b)-(i) and $\text{anchor}_3(u_{22}) = (u_{22}, u_{24})$ illustrates case (b)-(iii). Note that if there is only one edge (u, v) of Y_4^∞ incident to u in its cone i but v has two or more edges of Y_4^∞ incident to it in its cone $i + 2$, then $\text{anchor}_i(u)$ is not defined. For instance in Figure 5-(b), this is the case for $\text{anchor}_3(u_{14})$ and $\text{anchor}_1(u_7)$. It is always true, however, that if (u, v) is an edge lying in cone i of u then either $\text{anchor}_i(u)$ or $\text{anchor}_{i+2}(v)$ is defined. We use this to define a special type of path for every edge $(u, v) \in Y_4^\infty$:

Definition 3 Let (u, v) be an edge of Y_4^∞ lying in cone i of u such that $\text{anchor}_i(u) = (u, v')$ is defined. The *1-standard path* from u to v consists of edge (u, v') together with $\text{cpath}_u(v, v')$.

Since $\text{cpath}_u(v, v') \in Y_4^\infty$, there is a 1-standard path from u to v or from v to u for every edge $(u, v) \in Y_4^\infty$. If (u, v) is a dual edge in Y_4^∞ , there is a 1-standard path from v to u as well as one from u to v . The same is true if (u, v) is a mutually-single edge in cone i of u .

Lemma 4.1 *Let (u, v) be an edge of Y_4^∞ lying in cone i of u such that $\text{anchor}_i(u) = (u, v')$ is defined. Then (as illustrated in Figure 4-(b)):*

- (a) $d_2(u, v') \leq 2d_2(u, v)$,
- (b) the length of any edge in $\text{cpath}_u(v, v')$ is at most $\sqrt{2}d_2(u, v)$, and
- (c) the length of $\text{cpath}_u(v, v')$ is at most $(1 + \sqrt{2})d_2(u, v)$.

The length of the 1-standard path from u to v is thus at most $(3 + \sqrt{2})d_2(u, v)$.

Proof. We assume w.l.o.g. that $i = 0$. If $(u, v) = (u, v')$ the lemma trivially holds. Otherwise, let $(u, v_1), \dots, (u, v_k)$ be the fan of u in its cone 0 and let $(\overline{u}, \overline{v_l}) \in \overline{Y_4^\infty}$, $v' = v_{l'}$, and $v = v_r$ for some $l, l', r \in \{1, \dots, k\}$. We assume w.l.o.g. that $r > l'$. We assume that $r > l > l'$ as shown in Figure 4-(a) and Figure 4-(b). The proof for this case can be applied to prove the cases when $l \geq r$ or $l' \geq l$. By Definition 2, $\text{cpath}_u(v_{l'}, v_l)$ is a maximal uni-directional canonical path ending at v_l and, using Observation 3.1-(e), we have $d_2(u, v_{l'}) \leq d_1(u, v_{l'}) \leq d_1(u, v_l) \leq 2d_\infty(u, v_l)$. Since (u, v_l) is the shortest (with respect to the L_∞ -distance) edge in cone 0 of u , we have $2d_\infty(u, v_l) \leq 2d_\infty(u, v_r) \leq 2d_2(u, v_r)$. which proves part (a).

Any edge (v_s, v_{s+1}) , for $s \in \{l', \dots, l-1\}$, of $\text{cpath}_u(v_{l'}, v_l)$ must lie within rectangle $R(v_s, v_{s+1})$ and thus $d_2(v_s, v_{s+1}) \leq \sqrt{2}d_\infty(v_s, v_{s+1})$. Since $\text{cpath}_u(v_{l'}, v_l)$ is a maximal uni-directional canonical path ending at v_l and using Observation 3.1-(e), we have that $d_\infty(v_s, v_{s+1}) = d_y(v_s, v_{s+1})$ which in turn is at most $d_y(v_{l'}, v_l) \leq d_y(u, v_l) \leq d_\infty(u, v_l) \leq d_\infty(u, v_r) \leq d_2(u, v_r)$. Thus $d_2(v_s, v_{s+1}) \leq \sqrt{2}d_2(u, v_r)$. Any edge (v_s, v_{s+1}) for $s \in \{l, \dots, r-1\}$ must lie within $R(v_r, v_l)$ which in turn is contained in $S_u^i(d_\infty(u, v_r))$. Therefore $d_2(v_s, v_{s+1}) \leq \sqrt{2}d_\infty(u, v_r) \leq \sqrt{2}d_2(u, v_r)$ which completes the proof of part (b).

Using an argument similar to the one we used in the previous paragraph, $d_2(v_s, v_{s+1}) \leq \sqrt{2}d_y(v_s, v_{s+1})$, and hence, the length of $\text{cpath}_u(v_{l'}, v_l)$ is at most $\sqrt{2}d_y(v_{l'}, v_l)$. The length of




Figure 4: (a) The 1-standard path from u to $v = v_r$ consists of $\text{anchor}_i(u) = (u, v' = v_{l'})$ (in red) and $\text{cpath}_u(v' = v_{l'}, v = v_r)$ (in blue). Illustrated is the case when $(\overline{u}, \overline{v_l}) \in \overline{Y_4^\infty}$ for $r > l > l'$; $\text{cpath}_u(v' = v_{l'}, v_l)$ is a uni-directional canonical path. (b) Illustration of Lemma 4.1.

$\text{cpath}_u(v_r, v_l)$ is at most $d_1(v_r, v_l) = d_x(v_r, v_l) + d_y(v_r, v_l)$. So the length of $\text{cpath}_u(v_l, v_r)$ is at most $d_x(v_r, v_l) + d_y(v_r, v_l) + \sqrt{2}d_y(v_l, v_r) \leq d_x(u, v_l) + \sqrt{2}d_y(u, v_r) \leq d_\infty(u, v_l) + \sqrt{2}d_\infty(u, v_r) \leq (1 + \sqrt{2})d_\infty(u, v_r) \leq (1 + \sqrt{2})d_2(u, v_r)$, which completes the proof. \square

The reader can verify that the graph obtained by taking the union of 1-standard paths defined over all edges in Y_4^∞ , is a $(3 + \sqrt{2})$ -spanner of Y_4^∞ of maximum degree at most 12. We omit the proof because we do not make use of this fact in the rest of the paper.

Definition 4 An anchor (u, v) chosen by u in its cone i is *strong* if $\text{anchor}_{i+2}(v) = (v, u)$ or if $\text{anchor}_{i+2}(v)$ is not defined; it is *weak* if $\text{anchor}_{i+2}(v) \neq (v, u)$ (see Figure 5-(b)).

Let (u, v) be an edge of Y_4^∞ lying in cone i of u . If (u, v) is mutually-single then (u, v) is a strong anchor. For edges that are not mutually-single, we make the following observations:

Observation 4.2 Let u be a node that has a fan $(u, v_1), \dots, (u, v_k)$, with $k \geq 2$, in its cone i . Then $\text{anchor}_i(u)$ is defined, and if $\text{anchor}_i(u) = (u, v_{l'})$ for some $l' \in \{1, \dots, k\}$ then:

(a) If $l' = 1$ then $(\overrightarrow{v_1, v_2}) \in \overrightarrow{Y_4^\infty}$. If $l' = k$ then $(\overrightarrow{v_k, v_{k-1}}) \in \overrightarrow{Y_4^\infty}$. If $1 < l' < k$ then $(\overrightarrow{v_{l'}, v_{l'-1}}) \in \overrightarrow{Y_4^\infty}$ and $(\overrightarrow{v_{l'}, v_{l'+1}}) \in \overrightarrow{Y_4^\infty}$.

(b) If $(u, v_{l'})$ is a weak anchor then it is a dual edge.

Proof. By Definition 2-(b), if u has a fan of size at least 2 in its cone i then $\text{anchor}_i(u)$ is defined. Part (a) follows from Observation 3.1-(e) and Definition 2-(b). For part (b), note that if $(u, v_{l'})$ is a weak anchor then $\text{anchor}_{i+2}(v_{l'})$ is defined and is not $(u, v_{l'})$ which means that there must be two or more edges of Y_4^∞ incident to v in its cone $i + 2$. Therefore $(u, v_{l'})$ is a dual edge. \square

By Definition 2 and Definition 4, if $\text{anchor}_i(w_1) = (w_1, w_2)$ is a weak anchor then $\text{anchor}_{i+2}(w_2)$ is defined, and $\text{anchor}_{i+2}(w_2)$ could be either weak or strong. This means that starting from any weak anchor (w_1, w_2) there is a well-defined path of weak anchors $\text{anchor}_i(w_1) = (w_1, w_2)$, $\text{anchor}_{i+2}(w_2) = (w_2, w_3)$, $\text{anchor}_i(w_3) = (w_3, w_4)$, ... that would end when a strong anchor is encountered. Furthermore, if $\text{anchor}_i(w_1) = (w_1, w_2)$ is a weak anchor then any other anchor incident to w_1 in its cone i would have to be weak. Since weak anchors are dual and there can only be two dual edges incident to w_1 in its cone i , one of which is (w_1, w_2) , there can be only one other weak anchor incident to w_1 in its cone i , say $\text{anchor}_{i+2}(w_0) = (w_0, w_1)$. By repeatedly applying Observation 3.1-(d) to every successive pairs of nodes on $\text{cpath}_{w_j}(w_{j+1}, w_{j-1})$, we note that w_{j+1} is always in cone $i + 3$ of w_{j-1} and so the path w_0, w_1, \dots cannot form a cycle. This means that we can partition all weak anchors into maximal paths that we define as follows:

Definition 5 A *weak anchor chain* is a path w_0, w_1, \dots, w_k of maximal length consisting, for some $i \in \{0, 1, 2, 3\}$, of weak anchors $\text{anchor}_i(w_0) = (w_0, w_1)$, $\text{anchor}_{i+2}(w_1) = (w_1, w_2)$, $\text{anchor}_i(w_2) = (w_2, w_3)$, $\text{anchor}_{i+2}(w_3) = (w_3, w_4)$, and so on until:

- if k is even, weak anchor $\text{anchor}_{i+2}(v_{k-1}) = (v_{k-1}, v_k)$ such that $\text{anchor}_i(v_k) = (v_k, w)$ is a strong anchor, or
- if k is odd, weak anchor $\text{anchor}_i(v_{k-1}) = (v_{k-1}, v_k)$ such that $\text{anchor}_{i+2}(v_k) = (v_k, w)$ is a strong anchor.

For instance, in Figure 5-(b) $u_{22}, u_{24}, u_{15}, u_5$ and u_{28}, u_{26}, u_{27} and u_{21}, u_5 are three weak anchor chains. We now select anchor edges that will actually be included in the spanner:

Definition 6 We designate all strong anchors as *selected*. Furthermore, for every weak anchor chain w_0, w_1, \dots, w_k :

- (a) For $l = k - 1, k - 3, \dots$, we designate anchor (w_{l-1}, w_l) (chosen by w_{l-1}) as *selected*.
- (b) If (w_0, w_1) is not selected, i.e. k is odd, then we designate (w_0, w_1) to be a *start-of-odd-chain anchor* (chosen by w_0).

In Figure 5-(c), the dashed (not dotted) edges are the selected weak anchors; start-of-odd-chain anchors include edges $(u_3, u_{15}), (u_2, u_{10}), (u_{21}, u_5), (u_{17}, u_4), (u_{22}, u_{24})$ and (u_{13}, u_1) . The following observations regarding anchors are easy to check and the proofs are left to the reader:

Observation 4.3 For every node u and cone i of u :

- (a) There is at most one selected anchor incident to u in its cone i , whether the anchor is chosen by u or not.
- (b) If (u, v) is a start-of-odd-chain anchor chosen by u in its cone i then there is no selected anchor incident to u in its cone i .
- (c) If (u, v) is an anchor chosen by u in its cone i that is not selected, then there is a selected anchor chosen by v in its cone $i + 2$.
- (d) If (u, v) is an anchor chosen by u in its cone i that is not selected and that is not a start-of-odd-chain anchor, then there is another, selected anchor incident to u in its cone i (chosen by a node other than u).

We now define a new type of standard path that makes use of selected anchors only:

Definition 7 Let (u, v) be an edge of Y_4^∞ lying in cone i of u such that $\mathbf{anchor}_i(u) = (u, v')$ is defined. The 2-standard path from u to v is:

- The 1-standard path from u to v , if $\mathbf{anchor}_i(u)$ is selected.
- The path $\mathbf{cpath}_u(v, v')$ together with the 1-standard path from v' to u , if $\mathbf{anchor}_i(u)$ is not selected.

By Observation 4.3-(c), if $\mathbf{anchor}_i(u) = (u, v')$ is defined but not selected then $\mathbf{anchor}_{i+2}(v')$ is defined and selected. The 1-standard path from v' to u is thus well-defined and hence a 2-standard path from u to v is well-defined. By applying Lemma 4.1 twice, the length of the 2-standard path from u to v is at most $(3 + \sqrt{2})^2 d_2(u, v)$. In Figure 5-(a), the 2-standard path from u_{22} to u_{23} is path $u_{22}, u_{15}, u_{24}, u_{18}, u_{23}$. Note that non-selected anchors do not appear in 2-standard paths.

In the next section we will construct our first bounded degree spanner of Y_4^∞ . We will show that it contains short and highly structured paths between the endpoints of every edge in Y_4^∞ . We define the structure of these paths now:

Definition 8 Let H be a subgraph of Y_4^∞ that includes all selected edges, let (u, v) be an edge of Y_4^∞ lying in cone i of u such that $\mathbf{anchor}_i(u)$ is defined, and let p be the 2-standard path from u to v . For every $d \geq 1$, the $2d$ -standard pre-path in H from u to v consists of:

- (a) all edges on p that are in H
- (b) and, if $d > 1$, the $2(d-1)$ -standard pre-path in H from w to w' or from w' to w for every canonical edge (w, w') on p that is not in H .

When the $2d$ -standard pre-path in H is a path from u to v we call it the $2d$ -standard path or, more simply, the *standard path* when the value of d is understood from the context.

Because there is a 2-standard path for every $(u, v) \in Y_4^\infty$ and because H includes selected anchors, $2d$ -standard pre-paths in H are well-defined for all $(u, v) \in Y_4^\infty$. When a $2d$ -standard pre-path in H from u to v is a path, its length can be bounded by $(3 + \sqrt{2})^{2d} d_2(u, v)$ by applying Lemma 4.1 recursively.

5 A spanner of maximum degree at most 8

We now construct H_8 , our first bounded degree spanner of Y_4^∞ . It consists of all selected anchors and a subset of the uni-directional canonical edges of Y_4^∞ . We choose the edges in H_8 as follows:

Step 2 *We choose all the selected anchors. Then, for every node u and cone i of u , if $(u, v_1), \dots, (u, v_k)$, with $k \geq 2$, is the fan u in its cone i , we choose all the uni-directional edges on $\text{cpath}_u(v_1, v_k)$ except for the following two cases:*

- (a) $(v_2, v_1) \notin H_8$ iff (v_1, u) is a dual edge but not a start-of-odd-chain anchor chosen by v_1 and edge (v_2, v_1) is a non-anchor, uni-directional edge such that $(\overrightarrow{v_2, v_1}) \in \overrightarrow{Y_4^\infty}$ but $(\overrightarrow{v_1, v_2}) \notin \overrightarrow{Y_4^\infty}$.
- (b) $(v_{k-1}, v_k) \notin H_8$ iff (v_k, u) is a dual edge but not a start-of-odd-chain anchor chosen by v_k and edge (v_{k-1}, v_k) is a non-anchor, uni-directional edge such that $(\overrightarrow{v_{k-1}, v_k}) \in \overrightarrow{Y_4^\infty}$ but $(\overrightarrow{v_k, v_{k-1}}) \notin \overrightarrow{Y_4^\infty}$.

In order to facilitate the analysis of the maximum degree of the spanner, we devise a *charging scheme* that assigns each chosen spanner edge (u, v) to a cone of u and to a cone of v . A chosen edge (u, v) is charged to the cone of u containing it if it is a selected anchor or if it is a uni-directional canonical edge with $(\overrightarrow{u, v}) \in \overrightarrow{Y_4^\infty}$. By Observation 4.3-(a) and Observation 3.2-(b) at most one edge is charged to a cone in this way. If, however, edge (u, v) is a non-anchor, uni-directional canonical edge of some node w with $(\overrightarrow{v, u}) \in \overrightarrow{Y_4^\infty}$, i.e. its orientation in $\overrightarrow{Y_4^\infty}$ is incoming at u , it is charged to the cone of u that contains w . Therefore, to bound the degree of each node we only need to focus on cones that have one or more non-anchor, uni-directional, canonical, incoming (in $\overrightarrow{Y_4^\infty}$) edges charged to them.

We show in the following lemma that such cones have at most 2 edges charged to them, which will imply that H_8 has maximum degree at most 8:

Lemma 5.1 *Let u be a node with the fan $(u, v_1), \dots, (u, v_k)$, with $k \geq 2$, in its cone i and let $r \in \{1, \dots, k\}$. No more than 1 edge is charged to cone $i + 2$ of v_r except in the following cases when 2 edges are charged:*

- (a) $1 < r < k$, $(v_r, u) \notin H_8$ and both (v_{r-1}, v_r) and (v_r, v_{r+1}) are non-anchor uni-directional canonical edges in H_8 such that $(\overrightarrow{v_{r-1}, v_r}) \in \overrightarrow{Y_4^\infty}$ and $(\overrightarrow{v_{r+1}, v_r}) \in \overrightarrow{Y_4^\infty}$ (e.g., cone 3 of u_{14} in Figure 5-(c)).

- (b) $r = 1$, (v_1, u) is a non-anchor uni-directional canonical edge in H_8 such that $(\overrightarrow{v_1, u}) \in \overrightarrow{Y_4^\infty}$ and (v_2, v_1) is a non-anchor uni-directional canonical edge in H_8 such that $(\overrightarrow{v_2, v_1}) \in \overrightarrow{Y_4^\infty}$ (e.g., cone 3 of u_{21} Figure 5-(c)).
- (c) $r = k$, (v_k, u) is a non-anchor uni-directional canonical edge in H_8 such that $(\overrightarrow{v_k, u}) \in \overrightarrow{Y_4^\infty}$ and (v_{k-1}, v_k) is a non-anchor uni-directional canonical edge in H_8 such that $(\overrightarrow{v_{k-1}, v_k}) \in \overrightarrow{Y_4^\infty}$ (e.g., cone 0 of u_{19} Figure 5-(c)).

Proof. We consider first the case $1 < r < k$. If an incoming (in $\overrightarrow{Y_4^\infty}$) edge is charged to cone $i + 2$ of v_r then either $(\overrightarrow{v_r, v_{r-1}}) \notin \overrightarrow{Y_4^\infty}$ or $(\overrightarrow{v_r, v_{r+1}}) \notin \overrightarrow{Y_4^\infty}$. By Observation 4.2-(a), this means that $\text{anchor}_i(u) \neq (u, v_r)$. Furthermore, since (u, v_r) is a middle edge, by Observation 3.2-(c), (u, v_r) is not canonical and thus could not be chosen in Step 2. So $(u, v_r) \notin H_8$. Therefore, the maximum charge of 2 in cone $i + 2$ of v_r occurs when both (v_{r-1}, v_r) and (v_{r+1}, v_r) are charged to it which happens when they are both non-anchor uni-directional canonical edges with $(\overrightarrow{v_{r-1}, v_r}) \in \overrightarrow{Y_4^\infty}$, and $(\overrightarrow{v_{r+1}, v_r}) \in \overrightarrow{Y_4^\infty}$.

We consider case $r = 1$ next. Edge (v_2, v_1) is charged to cone $i + 2$ of v_1 only if (v_2, v_1) is a non-anchor, uni-directional edge, $(\overrightarrow{v_2, v_1}) \in \overrightarrow{Y_4^\infty}$, and either (v_1, u) is not dual or (v_1, u) is a start-of-odd-chain anchor chosen by v_1 .

In the second case, (v_1, u) is not a selected anchor and is thus not in H_8 . Furthermore, by Observation 4.3-(b), no other selected anchor is also charged to cone $i + 2$ of v_1 . Finally, there cannot be another start-of-odd-chain anchor out of v_i in its cone $i + 2$ because there is only one anchor out of a node in a cone. So cone $i + 2$ of v_1 could not be charged more than 1 in this case.

If (v_1, u) is not dual then it must be the only edge of Y_4^∞ in cone $i + 2$ of v_1 . Thus, for cone $i + 2$ to have a charge of 2, it must be that $(v_1, u) \in H_8$. Note that (v_1, u) cannot be an anchor chosen by v_1 because by Definition 2 v_1 chose no anchor. It also cannot be $\text{anchor}_i(u)$ because, using Observation 4.2-(a), that would violate the assumption that (v_2, v_1) is uni-directional and $(\overrightarrow{v_2, v_1}) \in \overrightarrow{Y_4^\infty}$. Therefore (v_1, u) must be a uni-directional canonical edge (with respect to some node w in cone $i + 3$ of u and v_1) and $(\overrightarrow{v_1, u}) \in \overrightarrow{Y_4^\infty}$.

The case $r = k$ follows by a symmetric argument to that in the case when $r = 1$. \square

We will show next that H_8 is a spanner of Y_4^∞ . In fact, we show something stronger:

Theorem 5.2 *Let (u, v) be an edge of Y_4^∞ lying in cone i of u . There is a 6-standard path in H_8 from u to v or from v to u of length at most $(3 + \sqrt{2})^6 \cdot d_2(u, v)$.*

We start by proving two special cases of the theorem:

Lemma 5.3 *Let (u, v) be a bi-directional canonical edge of Y_4^∞ lying in cone i of u such that $\text{anchor}_i(u)$ is defined. The 2-standard path from u to v is in H_8 and has length at most $(3 + \sqrt{2})^2 \cdot d_2(u, v)$.*

Proof. W.l.o.g. we assume that $i = 0$. The lemma clearly holds if $\text{anchor}_0(u) = (u, v)$ since (u, v) , being a canonical edge, would have to be a selected anchor and thus in H_8 . Therefore we can assume that $\text{anchor}_0(u) \neq (u, v)$ and, w.l.o.g., that $\text{anchor}_0(u)$ is clockwise from edge (u, v) in cone 0 of u . Then, since (u, v) is canonical, (u, v) must be the last edge in cone 0 of u . Let $(u, v_1), \dots, (u, v_k) = (u, v)$ be the fan of u in its cone 0 and let $\text{anchor}_0(u) = (u, v_{l'})$ for some $l' \in \{1, \dots, k - 1\}$.

Since (u, v_k) is bi-directional, $(\overrightarrow{u, v_k}) \in \overrightarrow{Y_4^\infty}$ and, by Definition 2, $\text{cpath}_u(v_l', v_k)$ is a maximal uni-directional canonical path ending at v_k . Since (u, v_k) is not dual, Step 2 ensures that $\text{cpath}_u(v_l', v_k = v)$ is a path in H_8 . Therefore, if $(u, v_l') \in H_8$ then the 1-standard (and thus 2-standard) path from u to v is in H_8 .

If $(u, v_l') \notin H_8$ then (u, v_l') must be a weak anchor and thus a dual edge. Since (u, v_l') lies clockwise from (u, v_k) within cone 0 of u , (u, v_l') must be the first edge in cone 0 of u , i.e. $l' = 1$. Any other anchor incident to u in its cone 0 would have to be a weak anchor and thus a dual edge (Observation 4.2-(b)). Since the first edge in the cone is (u, v_1) and the last is (u, v_k) , there cannot be another anchor incident to u in its cone 0. So (u, v_1) is a start-of-odd-chain anchor chosen by u . By Observation 4.3-(c), $\text{anchor}_2(v_1) \in H_8$ and $\text{anchor}_2(v_1) \neq (v_1, u)$. Since at least two edges of Y_4^∞ are incident to v_1 in its cone 2, node v_1 has a fan $(v_1, u_1), \dots, (v_1, u_{k'})$ with $k' \geq 2$ in its cone 2. Let $\text{anchor}_2(v_1) = (v_1, u_{l'})$ for some $l' \in \{1, \dots, k'\}$. Since, by Observation 3.2-(a), (v_1, u) is the first edge in cone 2 of v_1 , $u = u_1$. Because $(\overrightarrow{v_1, u}) \in \overrightarrow{Y_4^\infty}$ (a consequence of the assumption that $(\overrightarrow{u, v_k}) = (\overrightarrow{u, v}) \in \overrightarrow{Y_4^\infty}$), by Definition 2 $\text{cpath}_{v_1}(u_{l'}, u_1 = u)$ is a maximal uni-directional canonical path ending at $u = u_1$. Step 2 ensures that all the edges on this path are in H_8 (in particular $(u_2, u_1 = u)$ is in because (u, v_1) is a start-of-odd-chain anchor chosen by u). Therefore, the 2-standard path from u to $v = v_r$ is in H_8 . The bound on the length of the path follows from Lemma 4.1. \square

Lemma 5.4 *Let (u, v) be a uni-directional canonical edge of Y_4^∞ in cone i of u such that $(\overrightarrow{v, u}) \in \overrightarrow{Y_4^\infty}$ and $\text{anchor}_i(u)$ is defined. The 4-standard pre-path in H_8 from u to v is a path in H_8 and has length at most $(3 + \sqrt{2})^4 \cdot d_2(u, v)$.*

Proof. If $\text{anchor}_i(u) = (u, v)$ then (u, v) must be a selected anchor (only anchors that are dual may not be selected) and thus in H_8 so the theorem holds trivially. So we can assume that $\text{anchor}_i(u) \neq (u, v)$ and that, w.l.o.g., it lies clockwise from (u, v) within cone i of u .

If $\text{anchor}_i(u) = (u, v') \in H_8$ then the 2-standard path from u to v consists of edge (u, v') and $\text{cpath}_u(v', v)$, say $v' = v_{l'}, v_{l'+1}, \dots, v_k = v$. Step 2 ensures that all the uni-directional canonical edges on this path are in H_8 (in particular, edge $(v_{k-1}, v_k = v)$ is in because $(u, v_k = v)$ is a canonical edge and thus not dual). By Lemma 5.3, for every bi-directional edge in $\text{cpath}_u(v', v)$ not in H_8 , a 2-standard path between its endpoints is in H_8 . Therefore the 4-standard pre-path in H_8 from u to v is a path from u to v .

If $\text{anchor}_i(u) = (u, v') \notin H_8$ then the 2-standard path from u to v consists of $\text{cpath}_u(v', v)$, say $v' = v_{l'}, v_{l'+1}, \dots, v_k = v$, $\text{anchor}_{i+2}(v') = (v', u')$, and $\text{cpath}_{v'}(u', u)$, say $u = u_1, u_2, \dots, u_{l'} = u'$. Furthermore, just as in the proof of Lemma 5.3, $(u, v' = v_{l'})$ must be a start-of-odd chain anchor chosen by u . Step 2 ensures that all the uni-directional canonical edges on $\text{cpath}_u(v', v)$ and $\text{cpath}_{v'}(u', u)$ are in H_8 ; in particular, edge $(v_{k-1}, v_k = v)$ is in because $(u, v_k = v)$ is a canonical edge and thus not dual and $(u_2, u_1 = u)$ is in because $\text{anchor}_i(u)$ is a start-of-odd-chain anchor. By Lemma 5.3, for every bi-directional edge in $\text{cpath}_u(v', v)$ or $\text{cpath}_{v'}(u', u)$ that is not in H_8 , a 2-standard path between its endpoints is in H_8 . Therefore the 4-standard pre-path in H_8 from u to v is a path; the bound on its length follows from Lemma 4.1. \square

Proof of Theorem 5.2. The theorem holds trivially if (u, v) is a selected anchor, so we assume otherwise. W.l.o.g. we assume that $\text{anchor}_i(u)$ is defined and that $\text{anchor}_i(u)$ is either clockwise from (u, v) within cone i of u or that $\text{anchor}_i(u) = (u, v)$.

If $\text{anchor}_i(u) = (u, v')$ is in H_8 then the 2-standard path from u to v consists of edge (u, v') and $\text{cpath}_u(v', v)$, say $v' = v_l, v_{l+1}, \dots, v_r = v$. Step 2 ensures that all uni-directional canonical edges on this path are in H_8 except for possibly edge $(v_{r-1}, v_r = v)$; if missing, this edge is a uni-directional canonical edge such that $(\overrightarrow{v_{r-1}, v_r}) \in \overrightarrow{Y_4^\infty}$. By Lemma 5.4, the 4-standard pre-path in H_8 from v_r to v_{r-1} is a path from u to v . For every bi-directional edge in $\text{cpath}_u(v', v)$ not in H_8 , a 2-standard path between its end points is in H_8 . Therefore the 6-standard pre-path in H_8 from u to v is a path in H_8 .

If $\text{anchor}_i(u) = (u, v')$ is not in H_8 then the 2-standard path from u to v consists of $\text{cpath}_u(v', v)$, say $v' = v_1, v_2, \dots, v_r = v$, $\text{anchor}_{i+2}(v') = (v', u')$, and $\text{cpath}_{v'}(u', u)$, say $u = u_1, u_2, \dots, u_{l'} = u'$. Step 2 ensures that all uni-directional canonical edges on $\text{cpath}_u(v', v)$ and $\text{cpath}_{v'}(u', u)$ are in H_8 except for possibly $(v_{r-1}, v_r = v)$ and $(u_2, u_1 = u)$. By Lemma 5.4, the 4-standard pre-paths in H_8 from v_r to v_{r-1} and from u_1 to u_2 are paths in H_8 . By Lemma 5.3, for every bi-directional edge in $\text{cpath}_u(v', v)$ or $\text{cpath}_{v'}(u, u')$ that is not in H_8 , a 2-standard path between the endpoints is in H_8 . Therefore, the 6-standard pre-path in H_8 from u to v is a path and the bound on its length follows from Lemma 4.1. \square

6 Reducing the maximum degree bound to 4

In order to reduce the degree of H_8 , we need to remove for every cone with a charge of 2 at least one edge of H_8 that contributes to the charge. Lemma 5.1 describes the three cases in which a cone receives a charge of 2. We name the pair of non-anchor uni-directional canonical edges of u in its cone i satisfying the condition in case (a) of Lemma 5.1 an *edge pair in cone i of u* . We also name the non-anchor uni-directional canonical edge of u satisfying case (b) the *duplicate first edge in cone i of u* and the non-anchor uni-directional canonical edge of u satisfying case (c) the *duplicate last edge in cone i of u* .

We have shown in Theorem 5.2 that there is a 6-standard, or simply standard, path in H_8 between the endpoints of every edge in Y_4^∞ . These paths together satisfy the following:

Observation 6.1 *For every node u with a fan $(u, v_1), \dots, (u, v_k)$, with $k \geq 2$, in its cone i and $\text{anchor}_i(u) = (u, v_{l'})$ for some $l' \in \{1, \dots, k\}$, the following hold for the set of all 6-standard paths in H_8 :*

- (a) *For every edge pair $(v_{r-1}, v_r), (v_r, v_{r+1})$ of u (where $r \in \{2, \dots, k-1\}$) if a standard path in H_8 contains both edges then they must appear consecutively in the path, and if a standard path contains just one of them then the standard path must be from u to v_r .*
- (b) *If (v_2, v_1) is a duplicate first edge (of u) then no standard path in H_8 can contain (v_2, v_1) other than the standard path from u to v_1 .*
- (c) *If (v_{k-1}, v_k) is a duplicate last edge (of u) then no standard path in H_8 can contain (v_{k-1}, v_k) other than the standard path from u to v_k .*

Proof. By Definition 8, if (v_r, v_{r+1}) , for some $r \in \{1, \dots, k-1\}$, is a non-anchor uni-directional canonical edge that appears on a $2d$ -standard path in H_8 , say path p , from some node u' to another node v' , then either 1) (v_r, v_{r+1}) belongs to the 2-standard path p_1 from u' to v' or 2) (v_r, v_{r+1}) belongs to the $2(d-1)$ -standard path in H_8 from w to w' for some canonical edge (w, w') in p_1 that is not in H_8 . In case 1), by definition of 2-standard paths, u' must be u and the subpath of

p_1 starting at u and ending with edge (v_r, v_{r+1}) is the 2-standard path from u to v_r , if $l' \geq r + 1$, or to v_{r+1} , if $l' \leq r$. By Definition 8, the 2d-standard path from u to v_r , if $l' \geq r + 1$, or v_{r+1} , if $l' \leq r$, is contained in p . In case 2), we apply recursion until we obtain that (v_r, v_{r+1}) belongs to the 2-standard path from w to w' for some canonical edge (w, w') not in H_8 . Using the above argument, w must be u , (v_r, v_{r+1}) must be contained in $\text{cpath}_u(v_{l'}, w')$, and (u, w') must be the first or last edge in cone i of u .

If (v_2, v_1) is a duplicate first edge of u then (v_1, u) is not an anchor and $l' > 1$. Hence the only 1-standard path that uses (v_2, v_1) is the 1-standard path from u to v_1 and the only 2-standard path that uses (v_2, v_1) is the 2-standard path from u to v_1 . Since (u, v_1) is in H_8 , (v_2, v_1) appears only in one standard path in H_8 , the one from u to v_1 . This proves part (b) and, by symmetry, (c). To prove part (a), suppose $(v_{r-1}, v_r), (v_r, v_{r+1})$, for some $r \in \{2, \dots, k-1\}$, is an edge pair of u . By Observation 4.2-(a), $r \neq l'$. If only (v_{r-1}, v_r) appears in standard path p then $l' < r$ and, since (u, v_r) is not canonical, p must be the standard path from u to v_r . Similarly, if only (v_{r+1}, v_r) appears in standard path p then $l' > r$ and p must also be the standard path from u to v_r . Finally, if both edges appear in standard path p then they must be, by Definition 8, consecutive edges in the path. \square

The following observation applies to H_8 but we find it useful to state it more generally:

Observation 6.2 *Let H be a subgraph of Y_4^∞ that includes all selected edges, let $(u, v_1), \dots, (u, v_k)$, with $k \geq 2$, be the fan of cone u in its cone i , and let $\text{anchor}_i(u) = (u, v_{l'})$ for some $l' \in \{1, \dots, k\}$. If (u, v) , where $v = v_1$ or $v = v_k$, is a non-anchor uni-directional canonical edge then (u, v) cannot appear on the 6-standard pre-path in H from u to v .*

Proof. We assume w.l.o.g. that $i = 0$ and $(u, v) = (u, v_k)$. In that case, (u, v_k) is a canonical edge of some node s in its cone 3 where s lies within cone 1 of u . Note that edge (u, v_k) and anchor edge $(u, v_{l'})$ lie in cone 0 of u and cones 2 of v_k and $v_{l'}$, respectively.

By applying Observation 3.1-(d) to cone 0 of u and, if $(u, v_{l'})$ is not selected and thus $l' = 1$, cone 2 of v_1 , all non-anchor canonical edges on the 2-standard path p from u to v_k lie in cones 1 or 3 of their endpoints. By repeating this for every non-anchor canonical edge on p , we deduce that all non-anchor canonical edges on the 4-standard pre-path in H from u to v_k but not on the 2-standard pre-path in H lie in cones 0 and 2 of their endpoints. We continue one more time to find that all non-anchor canonical edges on the 6-standard pre-path in H from u to v_k but not on the 4-standard pre-path in H lie in cones 1 and 3 of their endpoints. This implies that if (u, v_k) does not appear in the 4-standard pre-path in H from u to v_k , it will not appear in the 6-standard pre-path in H either. Therefore we only need to show that (u, v_k) does not appear in the 4-standard pre-path in H from u to v_k .

The non-anchor canonical edges on the 2-standard path from u to v_k are 1) canonical edges of u in its cone 0 and, if $(u, v_{l'})$ is not selected and thus $l' = 1$, 2) canonical edges of node v_1 in its cone 2 and that lie in cone 3 of u . Consider a non-anchor canonical edge (v_r, v_{r+1}) of u in its cone 0. Either $\text{anchor}_1(v_r)$ is defined and $\text{anchor}_1(v_r) = (v_r, w')$ or $\text{anchor}_3(v_{r+1})$ is defined and $\text{anchor}_3(v_{r+1}) = (v_{r+1}, w')$. Note that w' must lie in cone 0 of u . The non-anchor canonical edges on the 2-standard path between v_r and v_{r+1} are all canonical edges of nodes v_r and w' if $\text{anchor}_1(v_r)$ is defined or v_{r+1} and w' if $\text{anchor}_3(v_{r+1})$ is defined. Since s lies in cone 1 of u , v_r , v_{r+1} , and w' must be different from s . If we now consider a canonical edge (w_1, w_2) of v_1 in its cone 2, we can similarly show that non-anchor canonical edges on the 2-standard path between w_1 and w_2 are canonical edges of nodes lying in cone 3 of u and thus cannot be s .

This implies that (u, v_k) , a canonical edge of s , cannot appear on a 4-standard pre-path in H

from u to v_k . □

By Observation 6.1-(b), if (v_2, v_1) is a duplicate first edge in cone i of u then no standard path in H_8 other than the 4-standard path (by Lemma 5.4) from u to v_1 uses edge (v_2, v_1) . Furthermore, edge $(v_1, u) \in H_8$ by definition (of duplicate first edge). Therefore, *as long as we keep edge (v_1, u)* , we can remove (v_2, v_1) from H_8 without breaking any standard path other than the one from u to v_1 and without increasing the stretch factor bound from Theorem 5.2. By symmetry, a similar insight can be made about Observation 6.1-(c) and duplicate last edge (v_{k-1}, v_k) . To generalize the discussion that follows, we will call an edge a *duplicate edge of u* if it is a duplicate first or last edge in some cone of u .

If (v_2, v_1) is a duplicate first edge of u in its cone i then, by Observation 3.2-(b), (v_2, v_1) is the last edge in cone $i + 1$ of v_1 and (u, v_1) is a non-anchor uni-directional canonical edge in H_8 with $(\overrightarrow{v_1}, \vec{u}) \in \overline{Y_4^\infty}$. Then, either (v_1, u) is a duplicate last edge in cone $i + 1$ of some node w lying in cones $i - 1$ of u and of v_1 or it is not a duplicate edge at all. This insight, along with the symmetric one regarding (v_{k-1}, v_k) and (v_k, u) , motivates this definition:

Definition 9 A *chain of duplicate edges* in H_8 is a path w_1, \dots, w_k, w_{k+1} of maximal length in H_8 in which every edge (w_{l-1}, w_l) is a duplicate edge of w_{l+1} for every $l = 2, 3, \dots, k$ and (w_k, w_{k+1}) is not a duplicate edge. Edge (w_k, w_{k+1}) is referred to as the *end edge* of the chain and k is the length of the chain.

For instance in Figure 5-(c), edge (u_3, u_2) is the end edge of the chain of duplicate edges u_3, u_2, u_{12}, u_{10} .

Observation 6.3 *The following hold for chains of duplicate edges:*

- (a) *A chain of duplicate edges does not form a cycle.*
- (b) *Every duplicate edge belongs to exactly one chain of duplicate edges.*
- (c) *Every non-anchor uni-directional canonical edge in H_8 that is not a duplicate edge is the end edge of exactly one chain, possibly of length 1.*

Proof. Let (w_{l-1}, w_l) be a duplicate (w.l.o.g., first) edge of w_{l+1} in its cone i and let w_{l-2} be the predecessor of w_{l-1} in a chain of duplicate edges containing (w_{l-1}, w_l) . Note that w_{l-2} lies in cone i of w_{l-1} . On the other hand, w_l, w_{l+1} , and all other successors of w_{l-1} on the chain lie in cone $i + 2$ or $i + 3$ of w_{l-1} . This proves part (a). Because (w_{l-2}, w_{l-1}) and (w_l, w_{l+1}) must be first edges in cones i of w_{l-1} and w_{l+1} , respectively, the edges that is before or after (w_{l-1}, w_l) in a chain of duplicate edges are uniquely defined proving part (b). For the same reason, there can be only one chain whose end edge is a given uni-directional non-duplicate edge and part (c) follows. □

By Observation 6.3, we can partition all non-anchor uni-directional canonical edges of H_8 into chains of duplicate edges. To construct our final spanner we first remove every other edge in every chain as follows:

Step 3 *For every chain of duplicate edges w_1, w_2, \dots, w_{k+1} we remove from H_8 every other edge in the chain starting with (w_{k-1}, w_k) , i.e. $(w_{k-1}, w_k), (w_{k-3}, w_{k-2}), \dots$*

We further remove edge pairs $(v_{r-1}, v_r), (v_{r+1}, v_r)$ and replace them with a shortcut:

Step 4 For every node u , every cone i of u , and every edge pair $(v_{r-1}, v_r), (v_{r+1}, v_r)$ in cone i of u , we remove (v_{r-1}, v_r) and (v_r, v_{r+1}) from H_8 , we add a new (straight-line) edge between v_{r-1} and v_{r+1} , and charge edge (v_{r-1}, v_{r+1}) to the cones of v_{r-1} and v_{r+1} in which the edge lies. We call this edge a shortcut between v_{r-1} and v_{r+1} . We also call v_r a cut-off node with respect to u .

Let H_4 be the resulting graph. In Figure 5-(d), edges (u_8, u_{14}) and (u_1, u_{14}) have been removed during step 4. All other edges present in H_8 but not in H_4 have been removed during step 3. Moreover, the shortcut edge (u_8, u_1) has been added during step 4.

Before we prove the main theorem, we show that the following strengthening of Theorem 5.2 holds for subgraph H_6 of H_8 that is obtained after applying Step 3, but not Step 4, to H_8 :

Lemma 6.4 Let (u, v) be an edge of Y_4^∞ lying in cone i of u . There is a 6-standard path in H_6 from u to v or from v to u of length no more than $(3 + \sqrt{2})^6 \cdot d_2(u, v)$.

Proof. Every standard path in H_8 that contains no edge removed in Step 3 is in H_6 , so we only need to consider standard paths that do contain a removed edge, i.e. get broken. Let (v_2, v_1) be a duplicate first edge, of some node u , that is in H_8 but not in H_6 (the argument for a duplicate last edge is symmetric). By Observation 6.1-(b), the only standard path in H_8 that gets broken by the removal of (v_2, v_1) is the standard path in H_8 from u to v_1 . Since (u, v_1) is a uni-directional canonical edge in Y_4^∞ , by Lemma 5.4 this standard path must be a 4-standard path p in H_8 from u to v_1 . Note that the standard path in H_8 from v_1 to v_2 , which happens to be a 4-standard path as well by Lemma 5.4, does not get broken (Observation 6.2) and so it is also the 4-standard path in H_6 from v_1 to v_2 . Since (u, v_1) is a uni-directional canonical edge and since (v_1, v_2) belongs to the 2-standard path from u to v_1 , the path from u to v_1 consisting of the 4-standard path from u to v_1 in H_8 with edge (v_2, v_1) being replaced by the 4-standard path from v_1 to v_2 in H_8 is a 6-standard path in H_6 from u to v_1 . \square

Theorem 6.5 H_4 is plane spanner of the Euclidean graph of maximum degree at most 4 and stretch factor at most $\sqrt{4 + 2\sqrt{2}}(1 + \sqrt{2})^2(3 + \sqrt{2})^6$.

Proof. We first argue planarity, which could potentially be affected by shortcut edges since they are the only edges in H_4 not in Y_4^∞ . Suppose (v_{r-1}, v_{r+1}) is a shortcut in H_4 that was put in because edge pair $(v_{r-1}, v_r), (v_{r+1}, v_r)$ in cone i of u was removed from H_8 and v_r is a cut-off node of u . Note that this implies that $(\overrightarrow{v_{r-1}}, \overrightarrow{v_r}), (\overrightarrow{v_{r+1}}, \overrightarrow{v_r}) \in \overrightarrow{Y_4^\infty}$. The only edge of Y_4^∞ that (v_{r-1}, v_{r+1}) intersects is (v_r, u) . That edge is a middle edge and hence not canonical. Therefore the only way for (u, v_r) to be in H_4 is if it was the anchor chosen by u . By Observation 4.2-(a), that would contradict the orientation of (v_{r-1}, v_r) and (v_{r+1}, v_r) in $\overrightarrow{Y_4^\infty}$. Furthermore, because a shortcut is added only between nodes v_{r-1} and v_{r+1} such that (v_{r-1}, v_r) and (v_{r+1}, v_r) are uni-directional, $(\overrightarrow{v_{r-1}}, \overrightarrow{v_r}) \in \overrightarrow{Y_4^\infty}$, and $(\overrightarrow{v_{r+1}}, \overrightarrow{v_r}) \in \overrightarrow{Y_4^\infty}$, it is not possible for two shortcut edges to intersect.

To argue the degree bound we only need to consider the three cases of Lemma 5.1. In case (b), either (v_1, u) or (v_2, v_1) is no longer in H_4 and the charge in cone $i + 2$ is reduced from 2 to 1. The same is true for case (c). In case (a), the charge in cone $i + 2$ of v_r is reduced from 2 to 0. The added shortcut edge (v_{r-1}, v_{r+1}) replaces the removed edges (v_{r-1}, v_r) and (v_{r+1}, v_r) and does not change the charge in the cones of v_{r-1} and v_{r+1} containing the shortcut edge.

In order to prove the stretch factor bound, by Lemma 6.4 we only need to consider the standard paths in H_6 that are broken by the removal of edges in Step 4. Consider an edge pair

$(v_{r-1}, v_r), (v_{r+1}, v_r)$ of some node u that is in H_8 but not in H_4 . Both edges are uni-directional canonical edges and thus, by Observation 3.2-(b), they are canonical edges of just one node (u). Because $(\overrightarrow{v_{r-1}}, \overrightarrow{v_r}), (\overrightarrow{v_{r+1}}, \overrightarrow{v_r}) \in \overline{Y_4^\infty}$, neither edge can be a duplicate edge and therefore both are in H_6 as well. If a standard path in H_8 contains both edges, by Observation 6.1-(a) the edges must be consecutive in the standard path and therefore shortcut edge $(v_{r-1}, v_{r+1}) \in H_4$ may be used instead of the missing edge pair which actually shortens the path and so these paths are not broken. If a standard path in H_8 uses just one of (v_{r-1}, v_r) or (v_{r+1}, v_r) then by Observation 6.1-(a) the standard path must be the one from u to node v_r , a cut-off node with respect to u in H_4 . Therefore, to complete the proof of the theorem, we only need to show that a short path exists in H_4 between u and cut-off node v_r (with respect to u). We assume w.l.o.g. that the standard path in H_8 from u to v_r uses edge (v_{r-1}, v_r) .

By Observation 6.2, edge (v_{r-1}, v_r) cannot appear in the 6-standard path in H_6 from v_r to v_{r-1} so its removal in Step 4 does not break that path. Furthermore, since v_r is a cut-off node with respect to u then (u, v_r) must be a middle edge and therefore v_{r-1} cannot be a cut-off node of v_r (since (v_r, v_{r-1}) is canonical and thus not a middle edge). So, the 6-standard path in H_6 between v_r and v_{r-1} must still exist, with shortcuts replacing any edge pairs on the path, in H_4 . By Lemma 4.1-(b) and Lemma 6.4, the length of this path is at most $\sqrt{2}(3 + \sqrt{2})^6 \cdot d_2(u, v)$. Since v_{r-1} cannot be a cut-off node of u (because $(\overrightarrow{v_r}, \overrightarrow{v_{r-1}}) \notin \overline{Y_4^\infty}$), the 6-standard path in H_6 from u to v_{r-1} is also in H_4 . The length of this path is no more than $(3 + \sqrt{2})^6 \cdot d_2(u, v)$. Therefore there is a path in H_4 from u to v_r of length bounded by $(1 + \sqrt{2})(3 + \sqrt{2})^6 \cdot d_2(u, v_r)$ and so H_4 is a spanner with stretch factor at most $\sqrt{4 + 2\sqrt{2}}(1 + \sqrt{2})^2(3 + \sqrt{2})^6$ by Lemma 3.3. \square

7 Conclusion

The question that this paper addressed is: *What is the smallest maximum degree that can be achieved for plane spanners of complete Euclidean graphs?* The main result of this paper allows this question to be reformulated as follows: *Is it always possible to construct a maximum degree 3 plane spanner of complete Euclidean graphs?*

Given the L_∞ -Delaunay triangulation of the considered point set, the construction of H_4 can be done in linear time. The stretch factor bound from Theorem 6.5 is a very rough bound on the stretch factor of H_4 . Our main goal was to present a simpler proof showing that H_4 is spanner of maximum degree four. The bound can be easily improved with a more careful analysis (leading to a proof with more cases to consider). We have written a program that constructs spanner H_4 on a set of points (see <http://www.labri.fr/~bonichon/deg4>) and we have failed to obtain examples that give a spanner with stretch factor greater than 10. We believe that the real spanning ratio is much lower and that this construction may have practical applications.

There exists a distributed algorithm that can compute a plane spanner of maximum degree 6 (and we denote this spanner H'_6) with a constant number of rounds [BGHP10]. For the construction of H_4 , the number of necessary rounds is bounded (from below) by the length of longest weak anchor chain and the length of the longest duplicated chain, and such a chain can have a linear number of vertices.

In [BFvRV12a] it is shown that there exists a routing algorithm on H'_6 with a bound stretch factor. We leave open the question whether or not it is possible to obtain a similar result on H_4 . The construction of H'_6 has been extended to constraints graphs in [BFvRV12b]. Once again, we leave open the question that whether or not it is possible to obtain similar results on H_4 .

8 Figures

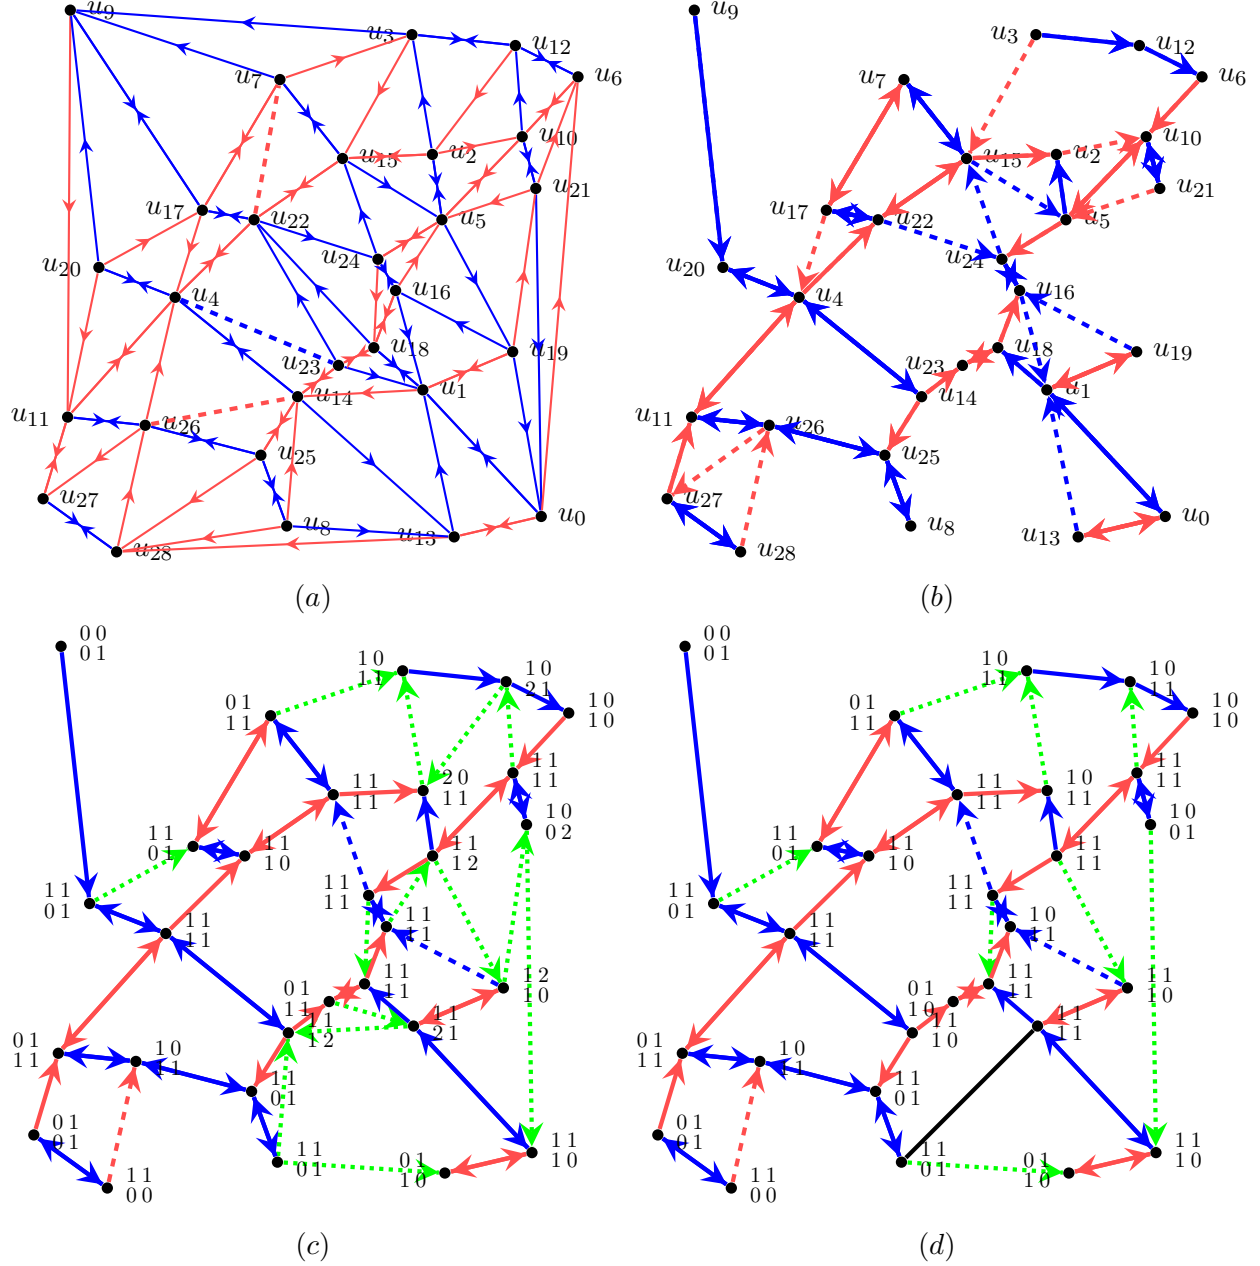


Figure 5: (a) The L_∞ -Delaunay triangulation of $P = \{u_0, u_1, \dots, u_{28}\}$. An edge lying in cone 0 of one endpoint and cone 2 of the other is colored red and an edge lying in cone 1 of one endpoint and cone 3 of the other is colored blue; the oriented edges belong to \bar{Y}_4^∞ while the dashed edges do not. (b) The anchor edges. An anchor (u_i, u_j) is shown to be oriented from u_i to u_j if it is chosen by u_i ; solid edges are strong anchors and dashed edges are weak anchors. (c) Graph H_8 . The label at a node shows the charge of each cone of that node; non-anchor uni-directional canonical edges are dotted green. (d) Graph H_4 . The undirected black edge is a shortcut edge.

References

- [BCCY12] P. Bose, P. Carmi, and L. Chaitman-Yerushalmi. On bounded degree plane strong geometric spanners. *J. Discrete Algorithms*, 15:16–31, 2012.
- [BDD⁺12] P. Bose, M. Damian, K. Douïeb, J. O’Rourke, B. Seamone, M. H. M. Smid, and S. Wuhler. $\pi/2$ -angle Yao graphs are spanners. *Int. J. Comput. Geometry Appl.*, 22(1):61–82, 2012.
- [BFvRV12a] P. Bose, R. Fagerberg, A. van Renssen, and S. Verdonschot. Competitive routing on a bounded-degree plane spanner. In *Proceedings of the 24th Canadian Conference on Computational Geometry (CCCG 2012)*, pages 299–304, 2012.
- [BFvRV12b] P. Bose, R. Fagerberg, A. van Renssen, and S. Verdonschot. On plane constrained bounded-degree spanners. In *Proceedings of the 10th Latin American Symposium on Theoretical Informatics (LATIN 2012)*, volume 7256 of *Lecture Notes in Computer Science*, pages 85–96, 2012.
- [BGHP10] N. Bonichon, C. Gavoille, N. Hanusse, and L. Perković. Plane spanners of maximum degree six. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 6198 of *Lecture Notes in Computer Science*, pages 19–30. Springer, 2010.
- [BGHP12] N. Bonichon, C. Gavoille, N. Hanusse, and L. Perković. The stretch factor of L_1 - and L_∞ -Delaunay triangulations. In *Proceedings of the 20th Annual European Symposium on Algorithms (ESA)*, volume 7501 of *Lecture Notes in Computer Science*, pages 205–216. Springer, 2012.
- [BGS05] P. Bose, J. Gudmundsson, and M. Smid. Constructing plane spanners of bounded degree and low weight. *Algorithmica*, 42(3-4):249–264, 2005.
- [BMSU01] P. Bose, P. Morin, I. Stojmenović, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.
- [BS13] P. Bose and M. Smid. On plane geometric spanners: A survey and open problems. *Computational Geometry*, 46(7):818–830, 2013.
- [BSX09] P. Bose, M. Smid, and D. Xu. Delaunay and diamond triangulations contain spanners of bounded degree. *Int. J. Comput. Geometry Appl.*, 19(2):119–140, 2009.
- [Che86] L. P. Chew. There is a planar graph almost as good as the complete graph. In *Proceedings of the Second Annual Symposium on Computational Geometry (SoCG)*, pages 169–177, 1986.
- [Che89] L. P. Chew. There are planar graphs almost as good as the complete graph. *Journal of Computer and System Sciences*, 39(2):205–219, 1989.
- [DFS90] D.P. Dobkin, S.J. Friedman, and K.J. Supowit. Delaunay graphs are almost as good as complete graphs. *Discrete & Computational Geometry*, 5(4):399–407, December 1990.
- [DH96] G. Das and P.J. Heffernan. Constructing degree-3 spanners with other sparseness properties. *Int. J. Found. Comput. Sci.*, 7(2):121–136, 1996.

- [KG92] J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry*, 7(1):13–28, 1992.
- [KP08] I. Kanj and L. Perković. On geometric spanners of Euclidean and unit disk graphs. In *Proceedings of the 25th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume hal-00231084, pages 409–420. HAL, 2008.
- [LW04] X.-Y. Li and Y. Wang. Efficient construction of low weight bounded degree planar spanner. *Int. J. Comput. Geometry Appl.*, 14(1-2):69–84, 2004.
- [Sal94] J. Salowe. Euclidean spanner graphs with degree four. *Discrete Applied Mathematics*, 54(1):55–66, 1994.
- [WL06] Y. Wang and X.-Y. Li. Localized construction of bounded degree and planar spanner for wireless ad hoc networks. *Mobile Networks and Applications*, 11(2):161–175, 2006.
- [Xia13] G. Xia. The stretch factor of the Delaunay triangulation is less than 1.998. *SIAM J. Comput.*, 42(4):1620–1659, 2013.